REMARKS/ARGUMENTS

In view of the foregoing amendments and the following remarks, the applicants respectfully submit that the pending claims are not anticipated under 35 U.S.C. § 102 and are not rendered obvious under 35 U.S.C. § 103. Accordingly, it is believed that this application is in condition for allowance. **If, however, the Examiner believes that there are any unresolved issues, or believes that some or all of the claims are not in condition for allowance, the applicants respectfully request that the Examiner contact the undersigned to schedule a telephone Examiner Interview before any further actions on the merits**.

The applicants will now address each of the issues raised in the outstanding Office Action.

Objections

Claim 5 is objected to because it recites acronyms. Claim 5 has been amended to replace the acronyms with there corresponding terms. (Claims 20 and 37 have been similarly amended.) Accordingly, this objection should be withdrawn.

Rejections under 35 U.S.C. § 102

Claim 29 stands rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Publication No. 2002/0176431A1 ("the Golla publication"). Since claim 29 has been canceled, this ground of rejection is rendered moot.

Rejections under 35 U.S.C. § 103

Claims 1-10, 16-23 and 34-37 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over the Golla publication in view of U.S. Publication No. 2003/0007498A1 ("the Angle publication"). The applicants respectfully request that the Examiner reconsider and withdraw this ground of rejection in view of the following.

Independent claims 1, 16 and 34 are not rendered obvious by the Golla and Angle publications because (i) one skilled in the art would not have been motivated to combine these publications as proposed by the Examiner, and (ii) because these publications neither teach, nor suggest, that each of the subschedulers requires more than one cell time slot to generate a match from its matching operation, but the subschedulers can collectively generate a match result for each output port in each cell time slot. These claims are reprinted below with this feature depicted in bold typeface:

> 1. For use with a switch having a first number of input ports, a second number of output ports, each of the input ports having the second number of virtual output queues, and a third number of subschedulers, each of the third number of subschedulers being able to arbitrate matching to each of the second number of output ports, a method for scheduling the dispatch of cells or packets stored in the virtual output queues, the method comprising:

a) for each of the virtual output queues, maintaining a first indicator for indicating whether the virtual output queue is storing a cell awaiting dispatch arbitration;
b) for each of the subschedulers, maintaining a second indicator F(i,j,k) for indicating whether the subscheduler is available or reserved; and
c) for each of the subschedulers, performing a matching operation, if it has been reserved, to match a cell buffered at a virtual output queue with its corresponding output port,

wherein each of the subschedulers requires more than one cell time slot to generate a match from its matching operation, and

wherein the subschedulers can collectively generate a match result for each output port in each cell time slot. [Emphasis added.]

16. For use with a switch including a first number of output ports, a second number of input ports, and the first number of virtual output queues associated with each of the second number of input ports, a dispatch scheduler comprising:

a) a third number of subschedulers;
b) a first indicator, associated with each of the virtual output queues, for indicating whether the virtual output queue is storing a cell awaiting dispatch arbitration; and
c) a second indicator, for each of the subschedulers, indicating

whether the subscheduler is available or reserved,

wherein each of the subschedulers is adapted to perform a matching operation, if it has been reserved, to match a cell buffered at a virtual output queue with its corresponding output port,

wherein each of the subschedulers requires more than one cell time slot to generate a match from its matching operation, and

wherein the subschedulers can collectively generate a match result for each output port in each cell time slot. [Emphasis added.]

34. For use with a switch having a first number of input ports, a second number of output ports, each of the input ports having the second number of virtual output queues, and a third number of subschedulers, each of the third number of subschedulers being able to arbitrate matching to each of the second number of output ports, a method for scheduling the dispatch of cells or packets stored in the virtual output queues, the method comprising for each of the subschedulers, performing a matching operation, if it has been reserved, to match a cell buffered at a virtual output queue with its corresponding output port,

wherein each of the subschedulers requires more than one cell time slot to generate a match from its matching operation,

wherein the subschedulers can collectively generate a match result for each output port in each cell time slot, and

fairness is maintained for best-effort traffic. [Emphasis added.]

This feature is discussed below. Then, the lack of suggestion in the art for the proposed combination is discussed.

The Examiner concedes that the Golla publication does not explicitly disclose that the subschedulers require more than one cell time slot to generate matching operations. (See, e.g., Paper No. 20041019, page 6.) To compensate for this admitted deficiency, the Examiner relies on the Angle publication. Specifically, the Examiner contends that page 11, paragraphs 115-116 of the Angle publication teach a scheduler performing a scheduling operation in more than one time slot, to generate select/grant to the request. The applicants respectfully disagree.

The Angle publication discusses a method of creating a combined schedule facilitated by pipelined staging of multicast and unicast scheduling. Specifically, the Angel publication discusses a method where multicast cells are scheduled for transmission first. Then, unicast cells are scheduled for transmission among the interfaces at a lower priority than the previously scheduled multicast cells, using only those interfaces that remain unmatched after completion of the multicast cell scheduling cycle. (See paragraph [0009].) As can be seen in Figure 11A of the Angle publication, multicast scheduling 1100 is performed every other cell timeslot and unicast scheduling 1110 is performed in every cell timeslot. A combined schedule 1120 is created by scheduling multicast cells to be dispatched in one or more cell timeslots later than the current cell timeslot in which both multicast cells and unicast cells are

present and are scheduled in parallel. (See FIG 11A, and paragraphs [0113]-[0116].) From the foregoing discussion, it is clear that the Angle publication does not disclose a method wherein ***each subscheduler requires more than one cell time slot to generate a match*** as recited by the claims. Thus, claims 1, 16 and 34 are not rendered obvious by the Golla and Angle publications for at least this reason. Since claims 2-10, 17-23, and 35-37 depend, either directly or indirectly, from claims 1, 16 and 34, respectively, these claims are similarly not rendered obvious by the Golla and Angle publications.

Further, one skilled in the art would not have been motivated to combine the Golla and Angle publications as proposed by the Examiner. The Angle publication concerns scheduling in a system including both unicast and multicast cells. However, the Golla publication does not disclose a scheduling system where both multicast and unicast scheduling may be performed. Therefore, there is no suggestion to one having ordinary skill in the art to modify the system of the Golla publication in view of the teachings of the Angle publication. Thus, claims 1-10, 16-23 and 34-37 are not rendered obvious by the Golla and Angle patents for at least this additional reason.

Further, with respect to independent claims 1 and 16, these claims are not rendered obvious by the Golla and Angle publications because these publications neither teach, nor suggest, a scheduler comprising a third number of subschedulers, each of which can arbitrate matching to each of the second number of output ports (by performing a matching operation) and each of which has a second indicator for indicating whether the subscheduler is available or reserved. The Examiner contends that the

Golla publication teaches these features. (See, e.g., Paper No. 20041019, pages 4 and 5.) The applicants respectfully disagree.

The Examiner contends that the system in the Golla publication uses a scheduler comprising a third number of subschedulers that are a combination of ingress servers/arbiters and egress servers/arbiters (citing FIGs. 1 and 5.), a second indicator (citing FIG. 12B, the combined matrix of Grants Matrix and scheduling matrix, citing FIG. 8, step 146.), for each of the subschedulers (citing FIG. 12B, the combined matrix contains/associates with each ingress/egress server.) indicating whether the subscheduler is reserved or available (citing page 8, paragraphs [0089]-[0090].). Note that the system of the Golla publication involves the interaction between the scheduler **and the output arbiters/servers** for proper cell scheduling. On the other hand, the claimed subschedulers can themselves arbitrate matching to each of the second number of output ports. As an example, in Figures 1 and 6 of the present application, the subschedulers only interact with an input port element for cell scheduling and do not need to interact with any output port element. In this exemplary embodiment, in the scheduler a third number of subschedulers, a second indicator for each of the subschedulers, and any other scheduler element all interact with input port element of the switch fabric. (See Figure 6; page 16, lines 30-31; and page 17, lines 1-23.)

Thus, claims 1 and 16 are not rendered obvious by the Golla and Angle publications for at least this additional reason. Since claims 2-10 and 17-23 depend, either directly or indirectly, from claims 1 and 16,

respectively, these claims are similarly not rendered obvious by the Golla and Angle publications.

Furthermore, independent claim 34 and dependent claims 4 and 19 recite that fairness for best effort traffic is maintained. The Examiner contends that the Golla publication discloses that fairness is maintained, citing paragraphs [0044], [0085], [0153] and [0156]. (See Paper No. 20041019, page 8.) Even assuming, arguendo, that this is true, the Examiner has elected to modify the Golla publication in view the Angle publication. As discussed above, the Angle publication teaches a method where multicast cells are given higher priority over unicast cells. (See, e.g., paragraph [0009].) Therefore fairness for best effort traffic cannot be maintained. Accordingly, these claims are not rendered obvious by the Golla and Angle publications for at least this additional reason. Since claims 35-37 depend from claim 34, they are similarly not rendered obvious by these publications.

Furthermore, claims 2, 3, 17, 18, 35 and 36 relate the number of cell time slots that each of the subschedulers need to generate a match to the number of subschedulers per input port. The Examiner (a) indicates that the Angle patent discloses using more than one time slot to perform scheduling/matching (See, e.g., Paper No. 20041019, page 7.), and (b) concludes that such a relationship is not patentable (See, e.g., Paper No. 20041019, page 8.). With respect to claims 2, 17 and 25, even assuming, arguendo, that the Angle patent discloses using more than one time slot to perform scheduling/matching, it does not disclose that such scheduling matching takes a number of time slots which is

the number of subschedulers. Moreover, the Examiner's conclusion that "If one has less number of server/scheduler/arbiter, it [scheduling] will require [a] lesser number of time slots, and [if] one has more number of server/scheduler/arbiter, it [scheduling] will require more [a greater] number of time slots," (Paper No. 20041019, page 8.) is untrue. Indeed, just the opposite is true. As shown in the exemplary embodiment of the present application depicted in Figure 8, the number of time slots required can match the number of subschedulers. This ensures that the subschedulers can collectively generate a match result for each output port in each cell time slot as claimed. This is not a mere design choice. Accordingly, claims 2, 3, 17, 18, 35 and 36 are not rendered obvious by the Golla and Angle publications for at least this additional reason.

Claims 11, 12, 24, 25, 30 and 31 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over the Golla publication in view of the Angle publication as applied to claims 1, 16 and 29, and further in view of U.S. Patent No. 6,141,323 ("the Rusu patent"). The applicants respectfully request that the Examiner reconsider and withdraw this ground of rejection in view of the following.

First, since claims 30 and 31 have been canceled, this ground of rejection is rendered moot with respect to these claims.

Second, the Examiner relies on the Rusu patent to teach a counter where the count is incremented upon learning that a new cell has arrived in the virtual output queue, and concludes that one skilled in the art

would have been motivated to combine this purported teaching into the proposed combination of the Golla and Angle publications. (See Paper No. 20041019, page 11.) However, even assuming, arguendo, that this is true, the purported teachings of the Russo patent fail to compensate for the deficiencies of the Golla and Angle publications with respect to claims 1, 16 and 34 described above. Accordingly, claims 11 and 12, and 24 and 25 are not rendered obvious at least for the reasons discussed above with respect to claims 1 and 16, respectively.

Claims 13-15, 26-28, 32 and 33 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over the Golla publication in view of the Angle publication as applied to claims 1, 16 and 29, and further in view of U.S. Patent No. 6,657,959 ("the Chong patent"). The applicants respectfully request that the Examiner reconsider and withdraw this ground of rejection in view of the following.

First, since claims 32 and 33 have been canceled, this ground of rejection is rendered moot with respect to these claims.

Second, the Examiner relies on the Chong patent as teaching a second indicator, for a subscheduler, set to indicate that the subscheduler is reserved if a first indicator indicates that a corresponding output is storing a connection/cell waiting to dispatch arbitration, and concludes that one skilled in the art would have been motivated to combine this purported teaching into the proposed combination of the Golla and Angle publications. (See Paper No. 20041019, page 13.)

However, even assuming, arguendo, that this is true, the purported teachings of the Chong patent fail to compensate for the deficiencies of the Golla and Angle publications with respect to claims 1, 16 and 34 described above. Accordingly, claims 13-15, and 26-28 are not rendered obvious at least for the reasons discussed above with respect to claims 1 and 16, respectively.

Conclusion

In view of the foregoing amendments and remarks, the applicants respectfully submit that the pending claims are in condition for allowance. Accordingly, the applicants request that the Examiner pass this application to issue.

Respectfully submitted,

January 28, 2005

John C. Pokotylo, Attorney
Reg. No. 36,242
Tel.: (732) 542-9070

CERTIFICATE OF MAILING under 37 C.F.R. 1.8(a)

I hereby certify that this correspondence is being deposited on **January 28, 2005** with the United States Postal Service as first class mail, with sufficient postage, in an envelope addressed to Mail Stop Amendment, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

John C. Pokotylo

36,242
Reg. No.